

06-16-00 17

FISH & RICHARDSON P.C.

4350 La Jolla Village Drive
Suite 500
San Diego, California
92122

Frederick P. Fish
1855-1930

June 14, 2000

Telephone
858 678-5070

Facsimile
858 678-5099

Web Site
www.fr.com

JC406 U.S. PTO
09/594352
06/14/00

W.K. Richardson
1859-1951

Attorney Docket No.: 10559/226001/P8791

Box Patent Application
Commissioner for Patents
Washington, DC 20231

Presented for filing is a new original patent application of:

Applicant: ERIK J. JOHNSON AND AARON R. KUNZE

Title: NETWORK ROUTING

Enclosed are the following papers, including those required to receive a filing date under 37 CFR 1.53(b):

	Pages
Specification	9
Claims	7
Abstract	1
Declaration	[To be Filed at a Later Date]
Drawing(s)	5

Enclosures:
— Postcard.

There are 30 total claims, 6 of which are independent.

Basic filing fee	\$0
Total claims in excess of 20 times \$18	\$0
Independent claims in excess of 3 times \$78	\$0
Fee for multiple dependent claims	\$0
Total filing fee:	\$0

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL558600700US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D C. 20231.

Date of Deposit June 14, 2000

Signature Derek W. Norwood

Typed or Printed Name of Person Signing Certificate
Derek W. Norwood



BOSTON

DALLAS

DELAWARE


NEW YORK

SAN DIEGO

SILICON VALLEY

TWIN CITIES

WASHINGTON, DC



FISH & RICHARDSON P.C.

Commissioner for Patents

June 14, 2000

Page 2

No filing fee is being paid at this time. Please apply any other required fees, **EXCEPT FOR THE FILING FEE**, to deposit account 06-1050, referencing the attorney document number shown above. A duplicate copy of this transmittal letter is attached.

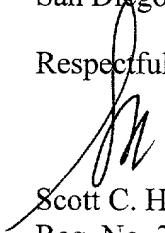
If this application is found to be incomplete, or if a telephone conference would otherwise be helpful, please call the undersigned at (858) 678-5070.

Kindly acknowledge receipt of this application by returning the enclosed postcard.

Please send all correspondence to:

SCOTT C. HARRIS
Fish & Richardson P.C.
Customer Number: 20985
4350 La Jolla Village Drive, Suite 500
San Diego, CA 92122

Respectfully submitted,



Scott C. Harris
Reg. No. 32,030
Enclosures
SCH/rpi
10038830.doc

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: NETWORK ROUTING

APPLICANT: ERIK J. JOHNSON AND AARON R. KUNZE

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL558600700US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

June 14, 2000
Date of Deposit


Signature

Derek W. Norwood
Typed or Printed Name of Person Signing Certificate

NETWORK ROUTING

BACKGROUND

This invention relates to routing data on a computer network.

5 A typical computer network is collection of interconnected computers. Data, such as text, images, sounds and videos, are typically communicated over the network from a source computer to a destination computer in pieces referred to as "packets." In most networks, packets usually are not sent directly from the source to the destination, but rather are routed by intermediate nodes known as "routers."

10 Generally, there are two types of packets that a router may receive. First, a general data packet contains data for communication between two computers on the network. The
15 router handles this type of data packet by forwarding the packet to the appropriate destination computer or next router as necessary.

20 Second, control information packets carry control information and are used by the router itself. Control information includes router-to-router communications and may also include administrative network commands. Changes of router configuration, such as management of communication ports and changing the internet protocol (IP) addresses of the

router communication ports, as well as changes of network topology, are examples of control information. Good network performance requires fast packet forwarding, but rapid exchange of control information is less critical to network performance.

To improve the speed of packet forwarding, many routers employ network processors, which are specialized hardware devices designed to perform certain networking operations quickly. A network processor operates separately from the router's general-purpose processor and usually isolates the operating system on the general-purpose processor from the forwarding operations. A router's network processor typically has more limited, dedicated resources than the router's general-purpose processor.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram a network having router interconnecting a number of computers.

Fig. 2 is a block diagram showing an interface between an operating system and a network processor.

Fig. 3 is a flow chart illustrating an embodiment of the invention.

Fig. 4 is a flow chart illustrating an embodiment of the invention.

Fig. 5 is a flow chart illustrating an embodiment of the invention.

DETAILED DESCRIPTION

5 The invention provides a non-invasive integration of network processors and operating systems. The invention is non-invasive in that the invention does not require any modifications to the operating system or to the network processor.

10 Although network processors have improved the speed of network routing, network processors still require the support of the operating system running on the general-purpose processor. In particular, many network processors are not equipped to handle control information, which must be handled
15 by the operating system on the general-purpose processor.

20 FIG. 1 is a high-level block diagram illustrating a system 2 in which a router 10 interconnects a number of computers 8 via network 4. Router 10 includes a number of forwarding interfaces 6 for receiving and transmitting packets over network 8.

 FIG. 2 is a block diagram depicting an example network router 10 and the internal communication paths between an operating system 12 and a network processor 32. The network

processor 32 includes N (an arbitrary number of) forwarding
interfaces 40, 42 and 44, connected by communication links 34,
36 and 38, respectively. By way of the forwarding interfaces
40, 42 and 44, the network processor 32 receives messages from
5 the network 4 or transmits messages over the network 4.

Forwarding interfaces 40, 42 and 44 represent connection ports
to other computers 8 in the network 4.

Separate from the network processor 32 is the router's
operating system 12, which typically runs on the router's
general-purpose processor (not shown). Drivers 20, 22 and 24,
also running on the general-purpose processor, act as
intermediaries between the operating system 12 and the network
processor 32. For each of the forwarding interfaces 40, 42
and 44, the operating system 12 instantiates a corresponding
driver. Thus, there are N forwarding interfaces 40, 42 and 44
and N corresponding drivers 20, 22 and 24.

Conventional operating systems include a number of
application program interfaces (API's), which are sets of
well-defined routines by which software applications interface
20 with the operating system. As such, conventional operating
systems typically include an API for network communication, by
which the operating system communicates with a driver for a
network interface card (NIC). In FIG. 1, however, no NIC's
are present in the router 10 because the network processor 32
25 handles network communications. Drivers 20, 22 and 24,

however, register with the operating system 12 as NIC device drivers and present a conventional interface to operating system 12. In this fashion, as described in detail below, router 10 supports the use of a dedicated network processor 32 yet provides a conventional interface to operating system 12.

When a driver registers with the operating system 12, the driver receives a routine from the operating system that the driver can call when the driver needs to deliver a control information packet to the operating system. The driver also provides a routine to the operating system 12 that the operating system can call when the operating system needs to deliver a control information packet to the network. The driver further informs the operating system 12 of the port or forwarding interface to which the driver corresponds. Through registration, communication links 14, 16 and 18 are established. The drivers 20, 22 and 24 present themselves to the operating system 12 as drivers for a network device, such as a driver for a NIC. Because the drivers 20, 22 and 24 appear to the operating system 12 as device drivers, the drivers 20, 22 and 24 can take advantage of API's for drivers already existing in the operating system 12.

Drivers 20, 22 and 24 also register with the network processor 32, and establish communication links 26, 28 and 30, respectively. Drivers 20, 22 and 24 present themselves to the network processor 32, not as drivers, but rather as forwarding

interfaces. To the network processor 32, drivers 20, 22 and 24 appear just like forwarding interfaces 40, 42 and 44. Control information to be sent from the network processor 32 to the operating system 12 is passed to a driver 20, 22 or 24, just as if the network processor 32 were forwarding the data along the network. The network processor 32 can use its own existing forwarding mechanisms to pass data to a driver 20, 22 or 24, which may then relay the data to the operating system 12.

Control information issuing from the operating system 12 can be relayed to the network processor 32 or to a forwarding interface 40, 42 or 44 by way of a driver 20, 22 or 24. The control information arrives in the network processor 36 like other data received from forwarding interfaces 40, 42 and 44. The network processor 32 then passes the information along the network by way of a forwarding interface 40, 42 or 44.

FIG. 2 illustrates an exemplary path 46 by which a general data packet is forwarded through the router 10. In FIG. 2, the general data packet is received at a first forwarding interface 40, which is then communicated 34 to the network processor 32. Because the general data packet includes no control information, the network processor 32 can ordinarily forward the data to a second forwarding interface 42, for transmission along the network. In a simple

forwarding operation, no intervention by the operating system 12 is required.

FIG. 2 also illustrates an exemplary path 48 of control information through the router 10. Control information may both be received by the operating system 12 from the network and sent from the operating system 12 by control information packets. When forwarding interface 44 receives a control information packet, the control data packet contained therein is communicated to the network processor 32. Control information entails the support of the operating system 12, so the network processor 32 can forward the control information to a driver 24, which then relays the control information to the operating system 12. Control information packets from the operating system 12 follow the reverse path, being communicated to the driver 24, which then forwards the control data packet to the network processor 32. The network processor 32 may then send the control information packet to a forwarding interface 44 for transmission along the network.

Protocols for driver API's vary from operating system to operating system. Similarly, protocols for forwarding interfaces vary from network processor to network processor. A driver may be created, however, to adapt an operating system with a driver API to a network processor supporting a forwarding interface. No modification of the operating system or network processor is needed.

FIG. 3 is a flow chart showing the steps for loading a driver. The driver itself is software, which may be stored in a computer system on any machine-readable medium such as a magnetic disk or optical drive, or may be stored within non-volatile memory such as read-only memory (ROM). An instance of a driver is to be created for each forwarding interface. This is accomplished by instantiating the driver on the general-purpose processor (50). The driver is registered with the operating system 12 so that the operating system's API will recognize the driver as a network device driver and so that communication with the operating system will be established (52). The driver is also registered with the network processor 32 so that the network processor will recognize the driver as a valid forwarding interface and so that communication with the network processor will be established (54).

FIG. 4 is a flow chart showing the steps for forwarding control information to the operating system 12. When the network processor 32 receives a data packet via a forwarding interface (60), the network processor 32 evaluates the data packet and determines whether it contains control information requiring the support of the operating system 12 (62). If no support of the operating system 12 is required, the network processor 32 routes the data packet to a forwarding interface (64) without involving the operating system. If the data

packet includes control information, however, the network processor 32 forwards the control information to a driver (66). The driver will then relay the control information to the operating system 12 (68). Relaying the control information can involve some reformatting of the information, i.e., putting the information in a form more useful to the operating system 12.

FIG. 5 is a flow chart showing the steps for passing control information from the operating system 12. The operating system 12 generates a control information packet (70), such as a router-to-router communication, and calls the routine that delivers the information to a driver (72). The driver in turn passes the control information packet to the network processor 32 (74), which sends the control information packet on the network 4 by way of the appropriate forwarding interface (76).

In a router with multiple forwarding interfaces, there are multiple drivers, one corresponding to each forwarding interface. The operating system selects for delivery of the control information the driver that corresponds to the forwarding interface that connects to the desired network port.

Other embodiments are within the scope of the following claims.

What is claimed is:

- 1 1. A method, comprising:
2 receiving a data packet transmitted over a network;
3 evaluating the data packet for the presence of control
4 information; and
5 if present, delivering the control information from a
6 network processor to an operating system by way of a non-
7 invasive driver.
8
9
10
11 2. The method of claim 1, wherein the operating system
12 provides an application programming interface (API) for
13 network communications, and further wherein the non-invasive
14 driver conforms to the API provided by the operating system.
15
16
17 3. The method of claim 1, wherein the non-invasive driver
18 appears to the operating system as a network interface device
19 driver.
20
21
22 4. The method of claim 1, wherein the non-invasive driver
23 appears to the network processor as a forwarding interface.
24
25
26 5. The method of claim 1, further comprising forwarding the
27 data packet on the network when control information is not
28 present in the data packet.

1 6. An article comprising a computer-readable medium which
2 stores computer-executable instructions for memory accessing,
3 the instructions causing a computer to:

4 receive a data packet transmitted over a network;

5 evaluate the data packet for the presence of control
6 information; and

7 if present, deliver the control information from a
8 network processor to an operating system by way of a non-
9 invasive driver.

10 7. The article of claim 6, wherein the operating system
11 provides an API for network communications, and further
12 wherein the non-invasive driver conforms to the API provided
13 by the operating system.

14 8. The article of claim 6, wherein the non-invasive driver
1 appears to the operating system as a network interface device
2 driver.

3 9. The article of claim 6, wherein the non-invasive driver
4 appears to the network processor as a forwarding interface.

5 10. The article of claim 6, the instructions further causing
6 a computer to:

3 when control information is not present in the data
4 packet, forward the data packet on the network.

1 11. A method, comprising:

2 instantiating a driver to communicate between an
3 operating system and a network processor;
4 registering the driver with the operating system; and
5 registering the driver with the network processor.

12. The method of claim 11, wherein the driver appears to the
operating system as a network interface device driver.

13. The method of claim 11, wherein the driver appears to the
network processor as a forwarding interface.

14. The method of claim 11, wherein the driver corresponds to
a forwarding interface between a network and the network
processor.

15. The method of claim 14, further comprising:

receiving a data packet at the forwarding interface, the
data packet including control information;
delivering the control information from the forwarding
interface to the driver; and

6 delivering the control information from the driver to the
7 operating system.

1 16. The method of claim 14, wherein the forwarding interface
2 is a first forwarding interface, the method further
3 comprising:

4 receiving a data packet at the first forwarding
5 interface, the data packet including data;

6 delivering the data of the data packet from the first
7 forwarding interface to a second forwarding interface; and
8 routing the data packet on the network.

1 17. The method of claim 16, wherein the data are delivered
2 from the first forwarding interface to the second forwarding
3 interface by the network processor.

1 18. The method of claim 14, further comprising:

2 generating control information in the operating system;
3 delivering the control information from the operating
4 system to the driver;

5 delivering the control information from the driver to the
6 network processor;

7 delivering the control information from the network
8 processor to the forwarding interface; and

9 routing the control information on the network.

1 19. An article comprising a computer-readable medium which
2 stores computer-executable instructions for memory accessing,
3 the instructions causing a computer to:

4 instantiate a driver to communicate between an operating
5 system and a network processor;

6 register the driver with the operating system; and

7 register the driver with the network processor.

20. The article of claim 19, wherein the driver appears to
the operating system as a network interface device driver.

21. The article of claim 19, wherein the driver appears to
the network processor as a forwarding interface.

22. The article of claim 19, wherein the driver corresponds
to a forwarding interface between a network and the network
processor.

23. The article of claim 22, the instructions further causing
a computer to:

receive a data packet at the forwarding interface, the
data packet including control information;

deliver the control information from the forwarding
interface to the driver; and

7 deliver the control information from the driver to the
8 operating system.

1 24. The article of claim 22, wherein the forwarding interface
2 is a first forwarding interface, the instructions further
3 causing a computer to:

4 receive a data packet at the first forwarding interface,
5 the data packet including data;

6 deliver the data of the data packet from the first
7 forwarding interface to a second forwarding interface; and
8 routing the data packet on the network.

1 25. The article of claim 24, wherein the data are delivered
2 from the first forwarding interface to the second forwarding
3 interface by the network processor.

1 26. The article of claim 22, the instructions further causing
2 a computer to:

3 generate control information in the operating system;

4 deliver the control information from the operating system
5 to the driver;

6 deliver the control information from the driver to the
7 network processor;

8 deliver the control information from the network
9 processor to the forwarding interface; and

10 route the control information on the network.

1 27. An apparatus comprising:

2 a general-purpose processor;

3 a driver executing on the general-purpose processor; and

4 a network processor configured to communicate data

5 packets to a set of forwarding interfaces,

6 wherein a subset of the forwarding interfaces is presented to

7 the network processor by the driver.

1 28. The apparatus of claim 27 further comprising an operating

2 system executing on the general-purpose processor, wherein the

3 operating system provides an API for network communications,

4 and further wherein the driver conforms to the API provided by

5 the operating system.

1 29. A driver comprising:

2 an operating system interface configured to appear to an

3 operating system as a network interface device driver; and

4 a network processor interface configured to appear to the

5 network processor as a forwarding interface.

1 30. The driver of claim 29, wherein the driver is configured

2 to take advantage of an API in the operating system.

NETWORK ROUTING

ABSTRACT OF THE DISCLOSURE

A method for employing a driver in a router to communicate between the router's operating system and network processor. The driver appears to the operating system as a device driver for a conventional network interface. The driver also registers with the network processor and appears itself as a forwarding interface. Control information received the network processor is forwarded to the operating system by way of the driver using the network processor's existing forwarding mechanisms, and control information from the operating system is sent to the network processor by way of the driver.


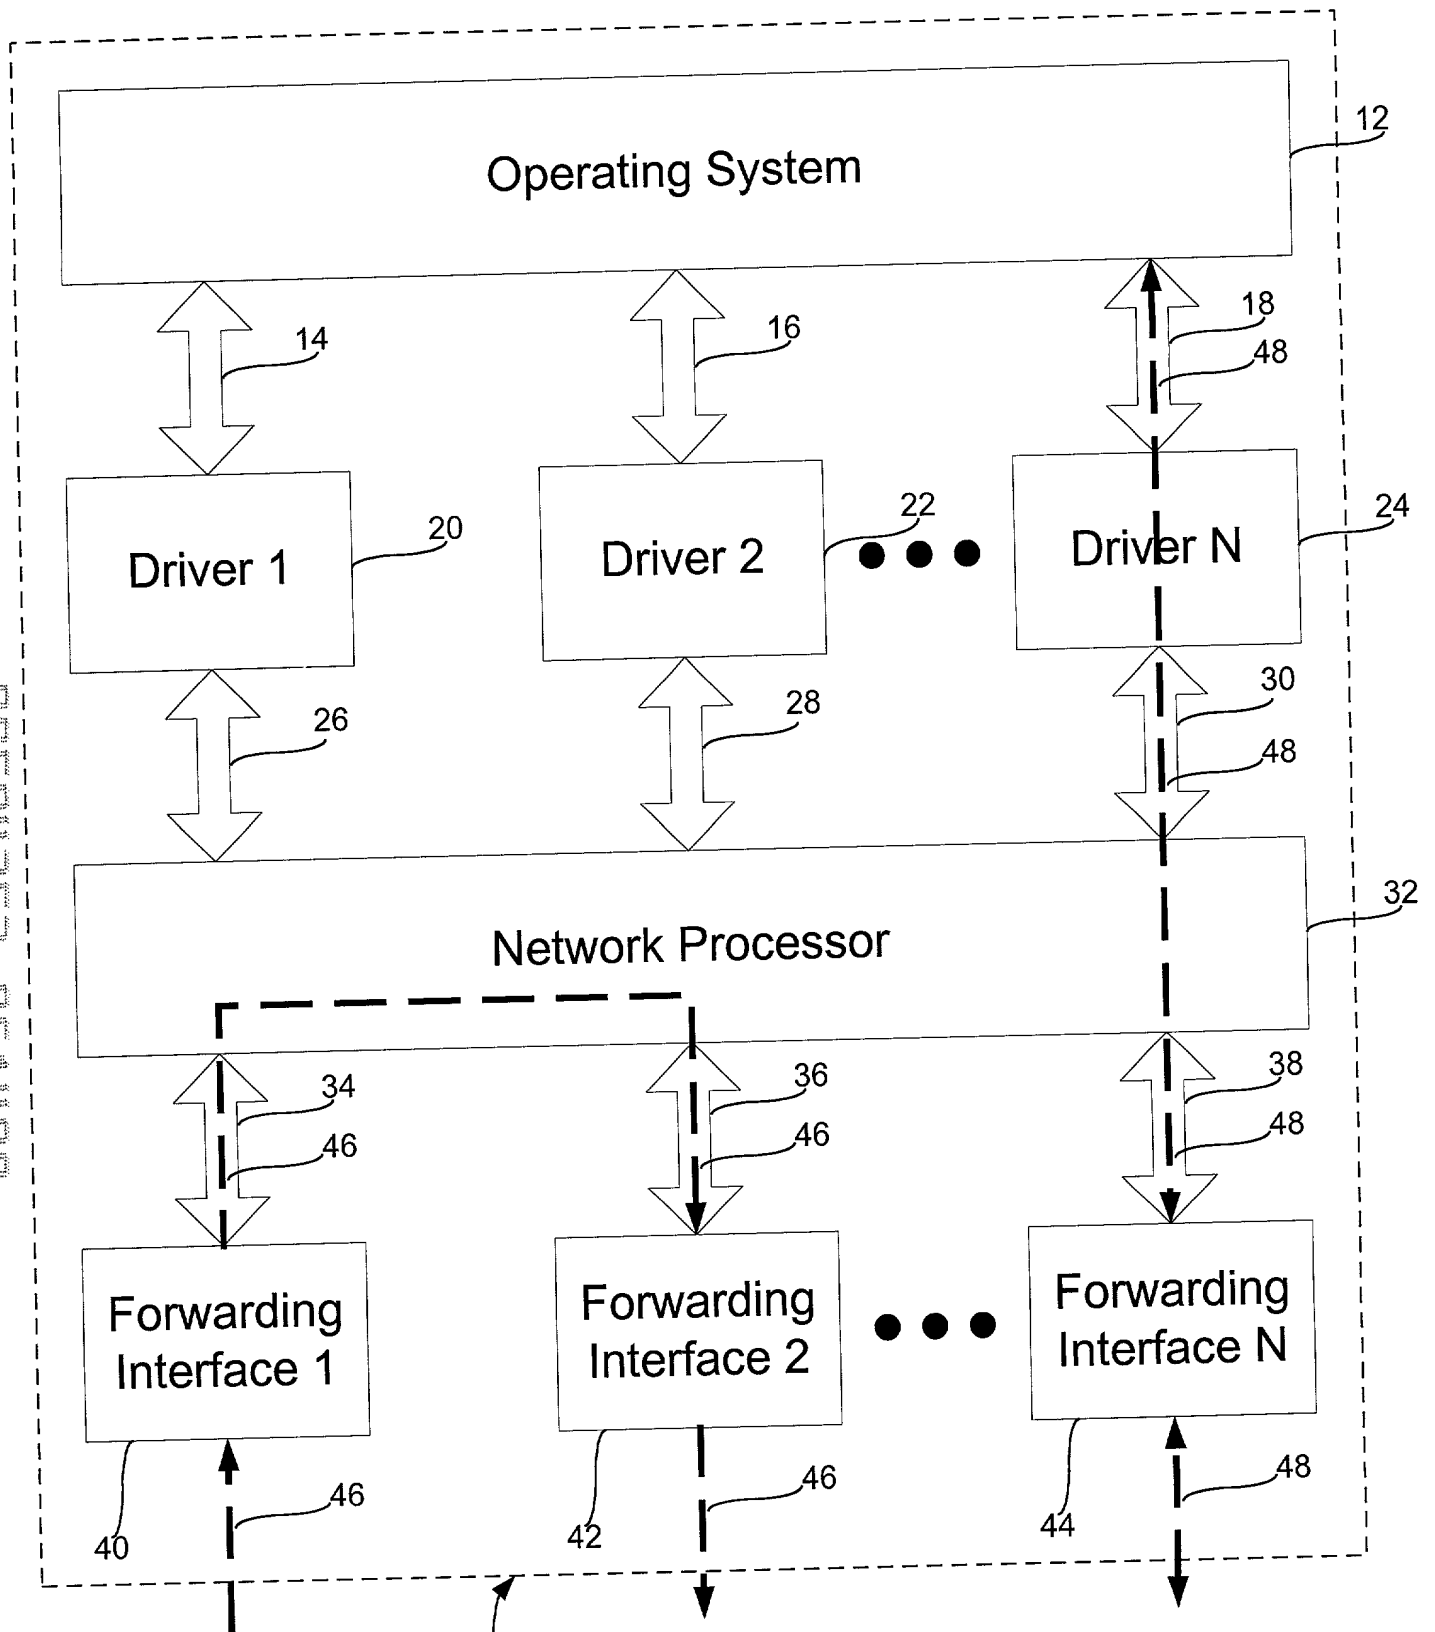


Fig. 1



10

Fig. 2

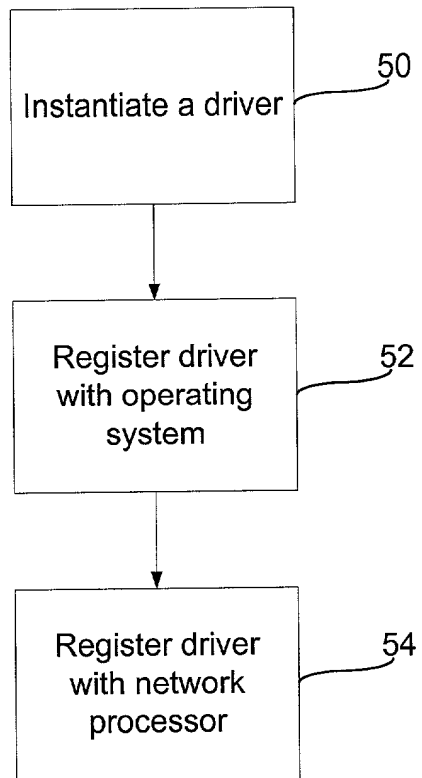


Fig. 3

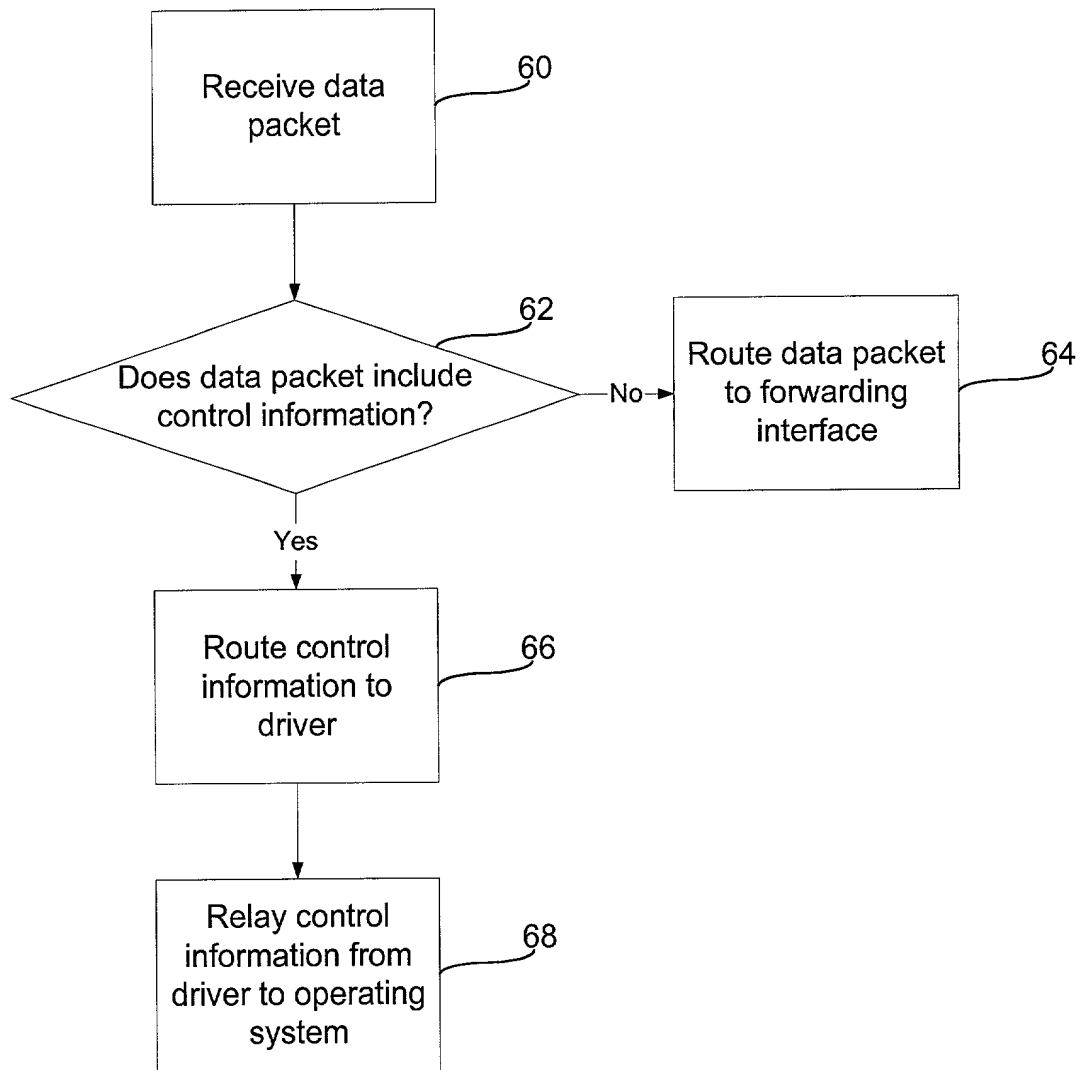


Fig. 4

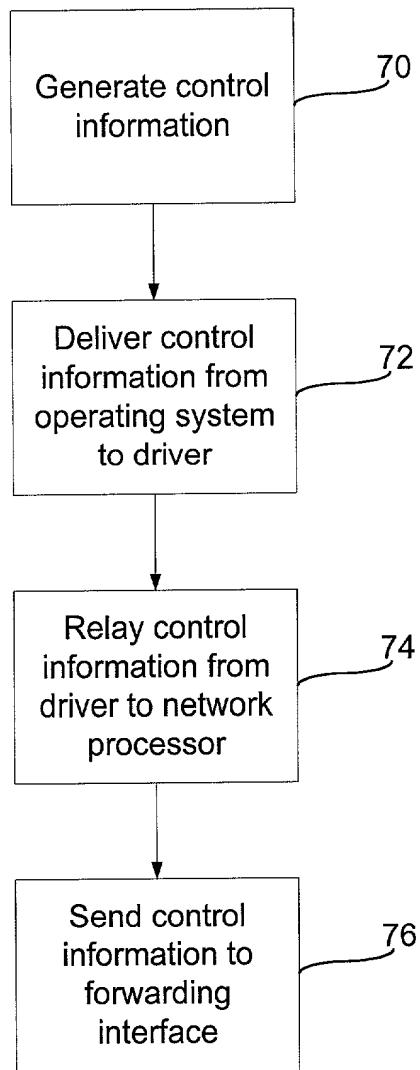


Fig. 5